

Keeping Researchers Connected

CERN IT simplifies HPC cluster connectivity with 10GbE iWARP from Intel

The European Organization for Nuclear Research, CERN, is the world's largest particle physics laboratory, located near Geneva, Switzerland. Its IT department is tasked with providing the information technology the laboratory needs to carry out its complex and innovative research both efficiently and effectively. This requires that the department not only develop leading research facilities, but also that it engineer and maintain a world-class infrastructure on which to run them.



CHALLENGES

- **Drive performance.** Boost research application results by enhancing cluster connectivity
- **Simplify management.** Reduce the complexity involved in managing network connections
- **Save money.** Create a more cost-effective connectivity model

SOLUTIONS

- **Simplified network fabric.** Replace the existing InfiniBand* fabric with 10GB Ethernet using Intel® Ethernet 10 Gigabit Server Adapters
- **High-performance compute.** Improve application performance with optimized compute in server clusters based on Intel® Xeon® processor 5500 series

IMPACT

- **Optimal performance.** Latency and application performance meet high-performance application requirements
- **Efficient user experience.** More flexible cluster connectivity makes resource allocation quicker and easier to manage, keeping the IT team and users happy
- **Reduced total cost of ownership.** The platform is simpler to maintain, enabling CERN IT to reduce TCO and focus on other areas

“With reduced maintenance complexity, we don’t need to invest as many resources in day-to-day upkeep of the cluster. Instead, we can focus on making sure it’s optimized to deliver the absolute best results for our research teams.”

Helge Meinhard,
Group leader, Platform and Engineering
Services Group, IT Department, CERN

Complex applications

One of CERN IT’s most important responsibilities is ensuring that the various engineering teams based at the facility have the tools they need to carry out calculations and simulations most effectively. It operates a wide variety of applications covering computational fluid dynamics (CFD), electronic data automation (EDA), structural and thermal analysis, and many other workloads.

A typical application is the CFD* software used for simulating heat distribution inside CERN’s particle accelerators and Large Hadron Collider* experiments. This is an essential process to ensure that some of the facility’s most challenging and high-profile experiments are carried out successfully.

These applications are often very complex, processing large volumes of data. They require a strong server platform to generate the best results — including excellent processing performance and optimized network connectivity. In terms of networking, latency and bandwidth are particularly crucial factors.

Many of the key applications, including STAR-CCM+* (CFD), GdfidL* (electromagnetic field calculations), ORBIT* (beam dynamics simulation), and Pelegant* (accelerator simulation), are run on a specialized server cluster, which previously used InfiniBand* adapters to provide connectivity.

The team managing the servers, headed by Helge Meinhard, group leader of the Platform and Engineering Services Group, wanted to replace the six-year-old cluster with newer, more powerful processor technology based on 10GB Ethernet Internet Wide-Area RDMA Protocol (iWARP) connectivity. “We wanted to simplify the connectivity we used, but also to improve the energy-efficiency and cost-effectiveness of our computing cluster,” explains Meinhard. Another important criterion was that the transition itself was as smooth and seamless as possible.



CERN IT switches to Intel® Ethernet 10 Gigabit Server Adapters to help enhance world-leading research capabilities

New server platforms

The original cluster was replaced with a new installation running on the Intel® Xeon® processor 5500 series to provide the level of compute performance required. To construct a cluster that was easier and more affordable to manage, while maximizing performance and minimizing latency, CERN chose to implement Intel® Ethernet 10 Gigabit Server Adapters as the fabric to connect the cluster.

A simplified networking solution

CERN has not typically used InfiniBand connectivity on a large scale, so the CERN IT networking team did not support it. The Intel Ethernet 10 Gigabit Server Adapter offered an easier-to-manage network that delivered the level of latency the engineering teams required. Meinhard comments: "We have chosen this solution as it gives us good performance based on standard protocols at reasonable cost and an acceptable power footprint."

After approximately six months of running the new cluster on the Intel Xeon processor 5500 series and 10GB Ethernet iWARP fabric, CERN IT was able to identify a significant ROI. The increased adaptability of the cluster, enabled by the new Ethernet connection, also made it more user-friendly, thereby meeting one of CERN IT's key objectives for the migration.

"We can run more parallel applications while giving users more support," Meinhard points out.

"With reduced maintenance complexity, we don't need to invest as many resources in day-to-day upkeep of the cluster," concludes Meinhard. "Instead, we can focus on making sure it's optimized to deliver the absolute best results for our research teams."

CERN IT supports complex applications for its engineering customers, and those applications demand high levels of performance and efficiency to provide accurate and timely results. RDMA (remote direct memory access) is a key element in the fabric that delivers the performance required, and has historically been available only on InfiniBand. Now iWARP, available from Intel® Ethernet, can provide this necessary RDMA capability over Ethernet. As CERN IT has found, using a high-performance Ethernet solution instead of InfiniBand enables great simplification in the fabric and higher ROI.

Performance: Data Intensive Computing: Support the most demanding business data processing, and computationally-intense graphics.

Find a solution that is right for your organization. Contact your Intel representative or visit the Reference Room at www.intel.com/references

Spotlight on CERN IT


CERN is a highly demanding computing environment. The CERN IT Department maintains extensive networks on both local and global scales. It pushes new technologies to their limits and provides a neutral ground for carrying out advanced R&D with various partners. Its services include large-scale computing and data storage, Internet exchange, telecom facilities, high speed networking, and database services. These services make a significant contribution towards helping CERN meet its primary objective of finding out what the universe is made of, and learn about the laws of nature.

The need for RDMA

High-performance computing (HPC) applications require extreme parallelism of the application across a cluster environment. With the advanced application performance and bandwidth available in the Intel Xeon processor 5500 series, high-performance networking is needed to efficiently connect the cluster.

HPC applications typically use Message Passing Interface (MPI) to deliver performance parallelism for the cluster application. MPI enables RDMA (remote direct memory access) to provide high-efficiency networking that reduces the CPU burden on applications.

RDMA solutions have long been available on InfiniBand. Not as widely known, iWARP provides the same RDMA mechanism over an Ethernet interface running TCP/IP. iWARP provides an Ethernet alternative to InfiniBand, while still providing the latency and performance required by these HPC applications. Being able to implement the performance RDMA protocol over a standard Ethernet fabric thus enables a new level of simplification in the fabric.



Copyright © 2011 Intel Corporation. All rights reserved. Intel, the Intel logo and Intel Xeon are trademarks or registered trademarks of Intel Corporation in the United States and other countries.

This document and the information given are for the convenience of Intel's customer base and are provided "AS IS" WITH NO WARRANTIES WHATSOEVER, EXPRESS OR IMPLIED, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT OF INTELLECTUAL PROPERTY RIGHTS. Receipt or possession of this document does not grant any license to any of the intellectual property described, displayed, or contained herein. Intel products are not intended for use in medical, life-saving, life-sustaining, critical control, or safety systems, or in nuclear facility applications.

Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance.

Intel may make changes to specifications, product descriptions and plans at any time, without notice.